TUM

TECHNISCHE UNIVERSITÄT MÜNCHEN
INSTITUT FÜR INFORMATIK

Complexity Results for Fork-Free Petri Nets

Ernst W. Mayr and Jeremias Weihmann

TUM-I1514

Complexity Results for Fork-Free Petri Nets

Ernst W. Mayr* Jeremias Weihmann*

January 26, 2015

Abstract

We investigate fork-free Petri nets (ff-PNs), which are those Petri nets for which each transition has at most one output place, connected by an arc with arbitrary multiplicity. We show that, for each reachable marking, there is a canonical firing sequence with nice properties leading to the marking. The existence of such sequences enables us to apply a mathematical framework for Petri nets which provides upper bounds for many classical problems. Using this framework and known lower bounds we show that the (zero-)reachability, boundedness, and covering problems of ff-PNs are **PSPACE**-complete. For the liveness problem of ff-PNs we obtain membership in **PSPACE**. Last, we show that the containment and equivalence problems of ff-PNs are **PSPACE**-hard and decidable in exponential space.

1 Introduction

Fork-free Petri nets (ff-PNs) are Petri nets, whose transitions have at most one outgoing arc (with arbitrary multiplicity). From the perspective of modeling, ff-PNs could be used to represent systems where different entities are transformed into many copies of the same end product. This class has already been considered by Holt and Commoner [2]. Lien [4] investigated termination properties of backward-concurrent-free Petri nets, a closely related subclass, where each transition has *exactly* one outgoing arc.Teruel and Silva [10, 11, 12] published a series of papers on equal conflict systems, a natural generalization of free-choice Petri nets containing the class of ff-PNs. Amer-Yahia et al. [1] presented approaches based on techniques of linear algebra to reason about properties of ff-PNs among other things. Morita and Watanabe [8] showed that the RecLFS problem of weighted state machines (WSMs), a subclass of ff-PNs where each transition additionally has at most one incoming arc, is **NP**-hard, even if restricted to WSMs with exactly two different arc multiplicities (the RecLFS problem asks, given a Petri net and a Parikh vector, if the Parikh vector is enabled in the Petri net). Taoka and Watanabe [9] investigated RecLFS for a subclass of WSMs with cactus structure. However, no improved upper bounds or even completeness-results for classical problems like RecLFS, reachability, boundedness, covering, containment or equivalence have been found so far. In this paper, we fill this gap.

We now describe our main tools. In [5], we showed (among other things) that the problems above are **PSPACE**-hard for WSMs, and therefore also for ff-PNs in general (we also refer to the journal version [7]). Furthermore, in Mayr and Weihmann [6], we provided a mathematical framework which can be used to obtain upper bounds for these problems if the Petri nets of the class under consideration allow for certain (canonical) firing sequences leading to reachable markings. In the same paper, they applied the framework to conservative Petri nets to obtain **PSPACE**-completeness for some of the problems above. This framework can also be used to obtain such results for join-free Petri nets (jf-PNs) which are those Petri nets for which each transition has at most one input place, connected with an arc with arbitrary multiplicity (i. e., they are the inverse nets of ff-PNs). Before the framework existed, we showed in [5], using similar ideas, that the RecLFS, (zero-) reachability, boundedness, and covering problems of jf-PNs are

*Institut für Informatik, Technische Universität München, D-85748 Garching, Germany, {mayr,weihmann}@in.tum.de

PSPACE-complete, and that the containment and equivalence problems are **PSPACE**-hard and decidable in doubly exponential space.

Usually, showing that the prerequisites of the framework are satisfied is the difficult part when using it to obtain upper bounds for computational problems of some class of interest. Accordingly, the major part of this paper is dedicated to showing that ff-PNs indeed have canonical firing sequences which are suited to apply the framework. It will turn out that our observation made in [5] that jf-PNs have such firing sequences is an important building block in our construction of canonical firing sequences for ff-PNs. The results we obtain are **PSPACE**-completeness for the (zero-)reachability, boundedness, and covering problems, **PSPACE**-membership for the liveness problem, and **PSPACE**-hardness and **EXPSPACE**-membership for the containment and equivalence problems.

This paper is organized as follows. In Section 2, we introduce concepts and notation that are used in later sections. In Section 3, we list the lower bounds of our problems of interest and observations about jf-PNs which we will use for our construction of canonical firing sequences in ff-PNs, both borrowed from [5]. In Section 4, we present the framework, which is borrowed from [6]. In Section 5, we discuss and describe, on a high level, our approach and the constructions of the following section (involving results for jf-PNs and the framework). In Section 6, we construct the canonical firing sequences of ff-PNs. In Section 7, we present the complexity results obtained by combining the known lower bounds with the upper bounds found by applying the framework. In Section 8, we conclude the paper with a short summary and outlook.

2 Preliminaries

2.1 Basic Notation and Encoding Scheme

Throughout this paper, we use the following notation to avoid confusion between elements of vectors or sequences and indexed elements of a set. We use $v_{[i]}$ in the few occasions we need to refer to the i-th element of a vector or a sequence v. The notation v_i is reserved for indexed elements of a set (e. g., $p_1, p_2, \ldots$). The set of all integers, all nonnegative integers, and all positive integers are denoted by $\mathbb{Z}$, $\mathbb{N}$, and $\mathbb{N}_{>0}$, respectively, while $[a, b] := \{a, a+1, \ldots, b\} \subsetneqq \mathbb{Z}$, and $[k] = [1, k] \subsetneqq \mathbb{N}_{>0}$. For two vectors $u, v \in \mathbb{Z}^k$, we write $u \geq v$ if $u_{[i]} \geq v_{[i]}$ for all $i \in [k]$, and $u > v$ if $u \geq v$ and $u_{[i]} > v_{[i]}$ for some $i \in [k]$. The maximum and minimum components of u are denoted by $\max(u) := \max_{i \in [k]} u_{[i]}$ and $\min(u) := \min_{i \in [k]} u_{[i]}$, respectively, where we define $\max(u) := \min(u) := 0$ if the dimension of the vector u is 0. Furthermore, let ld denote the function with $\mathrm{ld}(x) := \log_2(x)$ if $x > 0$, and $\mathrm{ld}(x) := 0$ otherwise. (Using this variation of the maximum and minimum functions, and of the logarithm lets us avoid special cases with Petri nets without places.) For two functions $f, g : X \to \mathbb{R}$, we write $f(a) \overset{P}{\leq} g(a)$ if there is a polynomial p such that $f(x) \leq p(g(x))$ for all $x \in X$. Note that $\overset{P}{\leq}$ is transitive.

Throughout this paper we use a succinct encoding scheme. Every number is encoded in binary representation. A vector of $\mathbb{N}^k$ is encoded as a k-tuple. A tuple is encoded as a tuple of the encodings of the particular components. For an object x, $\mathrm{size}(x)$ denotes the encoding size of x under this encoding scheme. The input size of a problem instance is the total size of the encodings of all entities that are declared as being "given" in the respective problem statement.

2.2 Petri Nets

A *Petri net* $N = (P, T, F)$ is a directed bipartite graph whose nodes consist of the *places* in P, *transitions* in T ($T \cap P = \emptyset$), and whose arcs and their *multiplicity* are given by the *flow function* $F : (P \times T) \cup (T \times P) \to \mathbb{N}$. Throughout this paper, n, m, and W will always refer to the number of places, transitions, and the largest arc multiplicity of the Petri net under consideration, resp. Usually, we assume an arbitrary but fixed order on P and T, respectively. With respect to this ordering of P, we can consider an n-dimensional vector v as a function of P, and, abusing notation,

write $v(p)$ for the entry of v corresponding to place p. Analogously, we write $v(t)$ in context of an m-dimensional vector v and a transition t.

A *marking* μ of N is a vector in $\mathbb{N}^n$. A pair (N, μ_0) such that μ_0 is a marking of N is called a *marked Petri net*, and μ_0 is called its *initial marking*. We will omit the term "marked" if the presence of a certain initial marking is clear from the context. The *inverse net* of N is obtained by inverting the direction of all arcs in N while keeping their multiplicities.

For a node $x \in P \cup T$, ${}^\bullet x := \{y \mid F(y,x) > 0\}$ ($x^\bullet := \{y \mid F(x,y) > 0\}$) is the *preset* (*postset*, resp.) of x. A Petri net is a *join-free* Petri net (jf-PN) if $|{}^\bullet t| \leq 1$ for all transitions t, a *fork-free* Petri net (ff-PN) if $|t^\bullet| \leq 1$ for all transitions t, and a *weighted state machine* (WSM) if it is a jf-PN and an ff-PN at the same time. For jf-PNs, we abuse notation by identifying ${}^\bullet t$ by its unique element (if this set is not empty). The transition t is *enabled at* μ or *enabled in* (N, μ) if $\mu(p) \geq F(p,t)$ for all $p \in P$. If t is enabled at μ, then it can be *fired*. By doing so, t produces the marking $\mu' = \mu + \Delta(t)$, where $\Delta(t)(p) = (F(t,p) - F(p,t))$ is the *displacement* of t at place p, and we write $\mu \xrightarrow{t} \mu'$.

Let $\sigma \in T^*$ be a transition sequence. The length of σ is denoted by $|\sigma|$, and ${}^\bullet\sigma = \bigcup_{i \in [|\sigma|]} {}^\bullet\sigma_{[i]}$ denotes its *preset*. We define $\sigma_{[i..j]} := \sigma_{[i]} \cdots \sigma_{[j]}$, $\sigma_{[..i]} := \sigma_{[1..i]}$, and $\sigma_{[i..]} := \sigma_{[i]}\sigma_{[i+1]} \cdots$. We write $\mu \xrightarrow{\sigma} \mu'$ if $\mu \xrightarrow{\sigma_{[1]}} \mu'' \xrightarrow{\sigma_{[2..]}} \mu'$ for some marking μ''. We then say that σ is enabled at μ, and leads from μ to μ', For the empty sequence ϵ, we define $\mu \xrightarrow{\epsilon} \mu$. If σ is enabled in (N, μ_0), then we call σ a *firing sequence*.

A marking μ is called *reachable* if $\mu_0 \xrightarrow{\sigma} \mu$ for some σ. The *reachability set* $\mathcal{R}(\mathcal{P})$ of a Petri net $\mathcal{P}$ consists of all markings reachable in $\mathcal{P}$. We say that a marking μ can be covered if there is a reachable marking $\mu' \geq \mu$. A Petri net $\mathcal{P}$ is called *bounded* if $\mathcal{R}(\mathcal{P})$ is finite. A Petri net $\mathcal{P}$ is called *live* if, for each transition t of $\mathcal{P}$ and each marking $\mu \in \mathcal{R}(\mathcal{P})$, there is a transition sequence enabled at μ that contains t.

The *displacement* of σ is defined by $\Delta(\sigma) := \sum_{i \in [|\sigma|]} \Delta(\sigma_{[i]})$. A *Parikh vector* Φ (also known as *firing count vector*) is simply an element of $\mathbb{N}^m$. The Parikh map $\Psi : T^* \to \mathbb{N}^m$ maps each transition sequence σ to its Parikh image $\Psi(\sigma)$ where $\Psi(\sigma)(t) = k$ for a transition t if t appears exactly k times in σ. The displacement of Φ is defined by $\Delta(\Phi) := \Delta(\sigma)$ for some σ with $\Psi(\sigma) = \Phi$. Abusing notation, we write $t \in \Phi$ if $\Phi(t) > 0$, and $t \in \sigma$ if $t \in \Psi(\sigma)$. The special Parikh vector $\Psi_{\text{first}}(\sigma)$ is a 0-1-vector such that, for all transitions t, $\Psi_{\text{first}}(\sigma)(t) = 1$ if and only if $t \in \sigma$ and ${}^\bullet t' \neq {}^\bullet t$ for all transitions $t' \in \sigma$ in front of the first occurrence of t in σ. For transition sequences σ and φ, the sequence that is obtained from σ by deleting the first $\min\{\Psi(\sigma)(t), \Psi(\varphi)(t)\}$ occurrences of every transition t, is denoted by $\sigma \dot{-} \varphi$.

A Parikh vector or a transition sequence with nonnegative displacement at all places is called *loop* since, if it can be fired at least once at some marking, the loop can immediately be fired again at the resulting marking. A loop with positive displacement at some place p is a *positive loop* (for p). A Parikh vector/transition sequence with nonpositive displacement at all places is a *nonpositive loop*. A nonpositive loop with negative displacement at some place p is a *negative loop* (for p). A loop with displacement 0 at all places is a *zero-loop*.

A Petri net is encoded as an enumeration of places $p_1, \ldots, p_n$ and transitions $t_1, \ldots, t_m$ followed by an enumeration of the arcs with their respective arc multiplicities.

In this paper, we consider the following classical computational problems for classes $\mathcal{C}$ of Petri nets.

- RecLFS: Given a PN $\mathcal{P} \in \mathcal{C}$ and a Parikh vector Φ, is Φ enabled in $\mathcal{P}$?
- Reachability: Given a PN $\mathcal{P} \in \mathcal{C}$ and a marking μ, is μ reachable in $\mathcal{P}$?
- Zero-reachability: Given a PN $\mathcal{P} \in \mathcal{C}$, is the empty marking reachable in $\mathcal{P}$?
- Covering: Given a PN $\mathcal{P} \in \mathcal{C}$ and a marking μ, is μ coverable in $\mathcal{P}$?
- Boundedness: Given a PN $\mathcal{P} \in \mathcal{C}$, is $\mathcal{P}$ bounded?
- Liveness: Given a PN $\mathcal{P} \in \mathcal{C}$, is $\mathcal{P}$ live?

- Containment: Given two PNs $\mathcal{P}$, $\mathcal{P}' \in \mathcal{C}$, is $\mathcal{R}(\mathcal{P}) \subseteq \mathcal{R}(\mathcal{P}')$?
- Equivalence: Given two PNs $\mathcal{P}$, $\mathcal{P}' \in \mathcal{C}$, is $\mathcal{R}(\mathcal{P}) = \mathcal{R}(\mathcal{P}')$?

3 Observations about Join-Free Petri Nets

In this section, we collect observations about jf-PNs which are needed later in Section 6. These are borrowed from [5]. We remark that we slightly changed the formulation of some lemmata compared to [5] to account for special cases like the net without places. For stronger versions of some of these statements we refer to the journal version [7] (which is to appear). A similar remark also applies to Section 4.

Lemma 3.1 ([5], Theorem 1)**.** *The RecLFS problem of general Petri nets is* **PSPACE***-complete, even if restricted to WSMs.*

Lemma 3.2 ([5], Lemma 2)**.** *Let σ be a firing sequence of a jf-PN (N, μ_0). If a transition $t \in \Psi_{first}(\sigma_{[i+1..]})$ is enabled at $\mu_0 + \Delta(\sigma_{[..i]})$, then $\sigma_{[..i]} \cdot t \cdot (\sigma_{[i+1..]} \dot{-} t)$ is a firing sequence.*

Lemma 3.3 ([5], Lemma 3)**.** *Let (P, T, F) be a jf-PN, σ a transition sequence, and μ, μ' markings with $\mu + \Delta(\sigma) = \mu'$ and $\mu(p)$, $\mu'(p) \geq W$ for all $p \in {}^\bullet\sigma$. Then, there is a permutation of σ enabled at μ (and leading to μ').*

Lemma 3.4 ([5], Lemma 5)**.** *Let $N = (P, T, F)$ be a Petri net. Then, there is a finite set $\mathcal{H}(N) = \{\Phi_1, \ldots, \Phi_k\} \subsetneq \mathbb{N}^m$ of loops of N such that each loop of $\mathcal{H}(N)$ consists of at most $(1 + (n+m)W)^{n+m}$ transitions, and such that, for each loop Φ of N, there are $a_1, \ldots, a_k \in \mathbb{N}$ with $\Phi = a_1\Phi_1 + \ldots + a_k\Phi_k$.*

Lemma 3.5 ([5], Lemma 6)**.** *There is a constant $c \in \mathbb{N}_{>0}$ such that, for each jf-PN $\mathcal{P} = (P, T, F, \mu_0)$ with $n > 0$ and each reachable marking μ of $\mathcal{P}$, there are $k \leq n \cdot \max(\mu)$ and transition sequences ξ, $\bar{\xi}$, α_1, ..., α_{k+1}, τ_1, ..., τ_k with the following properties, where $u = (nmW + \max(\mu_0) + 1)^{cn(n+m)}$:*

(a) $\xi = \alpha_1 \cdot \tau_1 \cdot \alpha_2 \cdot \tau_2 \cdots \tau_k \cdot \alpha_{k+1}$ is a firing sequence of length at most $(k+1)u$ leading from μ_0 to μ,

(b) $\bar{\xi} = \alpha_1 \cdot \alpha_2 \cdots \alpha_{k+1}$ is a firing sequence of length at most u, and

(c) each τ_i, $i \in [k]$, is a positive loop of length at most u enabled at some marking μ^* with $\max(\mu^*) \leq u$ and $\mu^* \leq \mu_0 + \Delta(\alpha_1 \cdot \alpha_2 \cdots \alpha_i)$.

Lemma 3.6 ([5], Theorems 2 and 3)**.** *The zero-reachability, the reachability, the boundedness, and the covering problems of jf-PNs are* **PSPACE***-complete, even if restricted to WSMs.*

Lemma 3.7 ([5], Theorem 4)**.** *The containment and the equivalence problems of jf-PNs are* **PSPACE***-hard, even if restricted to WSMs, and decidable in doubly exponential space.*

4 The Mathematical Framework

In this section, we describe the mathematical framework which we use to obtain complexity results for ff-PNs in Section 7. The central motives of this framework are *canonical firing sequences* ξ leading to reachable markings which result from inserting *short* loops $\tau_1, \ldots, \tau_k$ into *short backbone sequences* $\bar{\xi}$. It will turn out in Section 6 that, for ff-PNs, the backbones and loops have at most exponential length and all loops can be inserted at the same position at the backbone, which, using this framework, immediately implies the complexity results in Section 7. The framework is borrowed from [6].

Definition 4.1 (f-g-canonical class of Petri nets)**.** A class $\mathcal{C}$ of Petri nets is *f-g-canonical* for two monotonically increasing functions f, $g : \mathbb{N}^4 \to \mathbb{N}$ if, for each $\mathcal{P} = (P, T, F, \mu_0) \in \mathcal{C}$ with $n > 0$ and each marking μ reachable in $\mathcal{P}$, there are some $k \in [0, n(\max(\mu) + uW)]$ and transition sequences ξ, $\bar{\xi}$, α_1, …, α_{k+1}, τ_1, …, τ_k with the following properties, where $u = f(n, m, W, \max(\mu_0))$:

(a) $\xi = \alpha_1 \cdot \tau_1 \cdot \alpha_2 \cdot \tau_2 \cdots \alpha_k \cdot \tau_k \cdot \alpha_{k+1}$ is a firing sequence of length at most $(k+1)u$ leading from μ_0 to μ,

(b) $\bar{\xi} = \alpha_1 \cdot \alpha_2 \cdots \alpha_{k+1}$ is a firing sequence of length at most u,

(c) at most $g(n, m, W, \max(\mu_0))$ elements of $\{\alpha_1, \ldots, \alpha_{k+1}\}$ are nonempty sequences, and

(d) each τ_i, $i \in [k]$, is a positive loop of length at most u, enabled at some marking μ^* with $\max(\mu^*) \leq u$ and $\mu^* \leq \mu_0 + \Delta(\alpha_1 \cdot \alpha_2 \cdots \alpha_i)$.

An f-g-canonical class is

- *structurally f-g-canonical* if, for each $(N, \mu_0) \in \mathcal{C}$ and each marking μ of N, the Petri net (N, μ) is also in C, and
- *simple* if it can be determined in polynomial space if a given Petri net $\mathcal{P}$ belongs to $\mathcal{C}$, and if f and g are computable functions.

We remark that these canonical firing sequence leading to a reachable markings μ are not necessarily unique since there might be are other sequences satisfying the properties of Definition 4.1, and usually there are. However, we can assume w. l. o. g. that we always refer to unique sequences when considering the *canonical firing sequence*, for instance by considering the lexicographically smallest sequence satisfying this definition.

For classes of Petri nets satisfying this definition, several complexity results are known. By showing that a class is (structurally and simple) f-g-canonical, we obtain complexity results for problems involving this class for free.

Lemma 4.2 ([6], Theorem 2)**.** *Let $\mathcal{C}$ be a simple f-g-canonical class of Petri nets. Then, the reachability, and the covering problems are decidable in space polynomial in*

$$\mathrm{size}(\mathcal{P}) + \mathrm{size}(\mu) + n \operatorname{ld} f(n, m, W, \max(\mu_0)) + r,$$

and the boundedness problem is decidable in space polynomial in

$$\mathrm{size}(\mathcal{P}) + n \operatorname{ld} f(n, m, W, \max(\mu_0)) + r,$$

where r is the space needed to compute $f(n, m, W, \max(\mu_0))$.

Lemma 4.3 ([6], Theorem 4)**.** *Let $\mathcal{C}$ be a simple structurally f-g-canonical class of Petri nets. Then, the liveness problem of C is decidable in space polynomial in*

$$\mathrm{size}(\mathcal{P}) + n \operatorname{ld}(f(n, m, W, \max(\mu_0) + f(n, m, W, \max(\mu_0))W)) + r,$$

where r is the space needed to compute

$$f(n, m, W, \max(\mu_0) + f(n, m, W, \max(\mu_0))W).$$

Lemma 4.4 ([6], Theorem 5)**.** *Let $\mathcal{C}$ be a simple structurally f-g-canonical class of Petri nets. Then, for some polynomial p, the equivalence and containment problems of $\mathcal{C}$ are decidable in space*

$$p((K + 2u_1K)^{u_2 n} \cdot 2^{p(s + n \operatorname{ld}(u_3))} + r),$$

where

- *s is the encoding size of the input,*

- *n is the total number of places of both nets,*
- *m is the total number of transitions,*
- *W is the maximum of all arc multiplicities of both nets,*
- *K is the largest number of tokens appearing at some place at the initial markings,*
- *$u_1 = f(n, m, W, K)$, $u_2 = g(n, m, W, K)$, $u_3 = f(n, m, W, K + 2u_1W)$, and*
- *r is the time needed to compute u_1, u_2, and u_3.*

5 High Level Description

In the following, a sequence is called *long* if its length is at least exponential in the input size (where a non-specified exponential function separates between *long* and *short*).

A sufficient argument. Our goal is to show that the class of ff-PNs is f-2-canonical for an exponentially bounded function f and the constant function 2. Let μ be a marking reachable in an ff-PN $\mathcal{P}$, and σ be some firing sequence with $\mu_0 \xrightarrow{\sigma} \mu$. If μ has not substantially more tokens than μ_0, then we can use Lemma 3.5 (applied to the inverse net of $\mathcal{P}$ with μ as its initial marking) to argue that there is a *short* firing sequence leading from μ_0 to μ. This means that we are already finished for such a marking μ (in fact, the procedure that is described below also takes this case into consideration).

More interesting is the case where there is place p which contains substantially more tokens at μ compared to μ_0. Here, Lemma 3.5 cannot successfully be applied. Its application would possibly yield a long sequence ξ since the encoding size of μ (which is part of the input for the application of the lemma) is substantially larger than the size of the original input instance. This means, another approach must be found.

Consider a sub-procedure that takes a firing sequence σ with a big displacement as described above, leading from μ_0 to μ, and finds sequences α, τ, and β with the following properties.

(*i*) $\alpha \cdot \tau \cdot \beta$ is a firing sequence leading to μ,

(*ii*) $\alpha \cdot \beta$ is a firing sequence,

(*iii*) τ is a short positive loop.

Furthermore, we demand that applying the sub-procedure to σ and then iteratively to the residual sequences $\alpha \cdot \beta$ until $\alpha \cdot \beta$ has a small displacement yields short loops $\tau_1, \ldots, \tau_k$ with the property that

(*iv*) for each $i \in [k]$, the sequence $\alpha \cdot \tau_i$ is a firing sequence.

Note that if α or β are long, then, as discussed previously, we can use Lemma 3.5 to find short sequences α' and β' with respectively same displacements as α and β such that $\alpha' \cdot \beta'$ is a firing sequence.

The complete procedure of our interest consists of this iterative application of the sub-procedure described above, and of replacing of α and β by Parikh-equivalent sequences. Let α', τ_1, $\cdots$, τ_k, β' denote the sequences produced by this procedure.

If we had such a procedure, then this would immediately imply that the class of ff-PNs is f-2-canonical for some exponential function f. The reason is that we can simply consider each reachable marking μ, then take some firing sequence σ leading to μ, and apply the procedure to obtain $\xi = \alpha' \cdot \tau_1 \cdots \tau_k \cdot \beta'$. Here, the short sequences α' and β' are identified with the sequences α_1 and α_{k+1} of Definition 4.1, and $\alpha_2, \ldots, \alpha_k$ are empty sequences.

The sub-procedure is described in Lemma 6.1, although as a variation for jf-PNs because of reasons of technical and intuitive nature, which we explain later. The procedure is described in Corollary 6.3. The result that the class of ff-PNs is f-2-canonical can be found in Lemma 6.4.

A first idea for the sub-procedure and why it does not work. Let σ again denote a firing sequence from μ_0 to μ. We would like to argue that then σ contains a loop τ as a contiguous subsequence of σ such that the sequence which results from removing τ from σ is still a firing sequence. Indeed, if the number of tokens at p is doubly exponentially larger at μ compared to μ_0, then arguments involving the coverability tree (see Karp and Miller [3]) show that σ contains a positive loop τ such that $\sigma = \alpha \cdot \tau \cdot \beta$. However, it's not necessarily the case that $\alpha \cdot \beta$ is still firable because it may be that τ is necessary for β to be fired, i. e., property (ii) is not necessarily satisfied. Property (iii) is also not necessarily satisfies since this argument only provides a doubly exponential upper bound for the length of τ. Also, the upper bound on the length of σ (doubly exponential) and the prerequisites for this argument (doubly exponential difference in the token numbers) are rather weak. We need stronger bounds on these entities if we want to use our framework to obtain the desired upper bounds for the decision problems of our interest.

A better idea. We now describe a better approach on a very high level, that is, we address only the most important ideas without explaining the many technical details. Furthermore, this approach here is simplified compared to the actual one in the following chapters, and is therefore not feasible with respect to certain details. Hence, the reader is encouraged to only take the ideas, and not to try to understand every detail at this point.

Observe that if there is a place p such that $\mu(p) > \mu_0(p)$, then σ must contain a transition that increases the number of tokens at p. We want to argue that if $\mu(p) - \mu_0(p)$ is big, then there are enough such transitions within σ such that we can combine some of these together with a few other transitions, which balance out the decrease of tokens at places other than p, and obtain a loop τ. In other words, by permuting σ, we obtain α, τ, and β such that $\alpha \cdot \tau \cdot \beta$ is an enabled permutation of σ, and τ is a positive loop which increases the number of tokens at p.

The process of creating τ consists of iteratively choosing transitions from σ that either increase the number of tokens at p (because we want to have a positive loop for p) or increase the number of tokens at a place for which the current sequence τ decreases the number of tokens (because the displacement at each place must be nonnegative). When we pick a transition for τ, we want to do it in such a way that at each point $\alpha \cdot \tau \cdot \beta$ is enabled.

To this end, we observe that if α leads to some marking μ' then we can choose the most right transition within α that puts tokens to a place p with $\mu'(p) \geq W$, and push it to the last position in α. The resulting sequence is still enabled. In such a manner, we can let τ grow *in the backward direction*, i. e., we first choose its *last* transition by pushing it to the end, then the *second-last* transition by pushing it in front of the last transition, etc.

However, instead of the original ff-PN $\mathcal{P}$, we consider its inverse net $\mathcal{P}'$, which is a jf-PN, and describe a procedure where we let a *negative loop* grow in the *forward direction*. The reason why we do this with an jf-PN instead of an ff-PN is that this permutation procedure consists of many steps where we must ensure at each point that the permutation of the original sequence is also a firing sequence. This is easier for jf-PNs in an intuitive sense since we must only ensure that the unique input place of a transition under consideration contains enough tokens. By reversing the whole sequence later (which then is a firing sequence in $\mathcal{P}$), we obtain a positive loop.

To create τ, the main idea is the following: We start with an initial marking μ_0 and a firing sequence σ of $\mathcal{P}'$, leading to an end marking μ_σ, where there is a place p^* such that $\mu_0(p^*) - \mu_\sigma(p^*)$ is big. We let a prefix α' grow by iteratively picking transitions in an appropriate way from the shrinking suffix β (which is initialized by $\beta \leftarrow \sigma$). Let ν denote the marking reached by the current prefix α'. (Since α' changes during the procedure, ν also changes.)

The way we pick the next transition is as follows: We first define appropriate small intervals for each place $p \neq p^*$ such that the interval for p contains $\mu_0(p)$. Here we assume that the initial marking at places $p \neq p^*$ is not much larger than μ_σ in order to make the intervals small. (In this description, we leave out a technical preprocessing step which yields an alternative initial marking with this property, which is then considered instead of the real initial marking.) These intervals are defined in such a way that the permutation procedure can never push $\nu(p)$ below the lower boundary of the interval of p. The upper boundary for each place $p \neq p^*$ is at least $\mu_\sigma(p)$ in order

to make sure that we can pick a transition consuming tokens from p whenever its token number is above the upper boundary.

Each time $\nu(p)$ leaves the interval of p by crossing the upper boundary, we extend α' by choosing a transition which consumes tokens from p in order to push the number of tokens back into the interval. It could be the case that the transitions chosen in this manner will at some point constitute a positive loop since it can happen that they push tokens back and forth between big places and increase their total number of tokens. However, this is no issue at all since eventually the number of tokens will necessarily enter the interval again because the end marking at p is smaller than the the upper boundary of p. If, however, the token numbers at all places $p \neq p^*$ are within the respective intervals, we choose a transition that consumes tokens from p^*.

The important observation is that the number of different sub markings, restricted to the places $p \neq p^*$, where all token numbers are in their intervals, is small since the intervals are small. That means that the number of times we observe a new smallest marking w. r. t. the number of tokens at p^* where all token numbers at places $p \neq p^*$ are in their intervals, is also small. At some point, we repeat such a sub marking. The sequence τ between the first and the next time we observe this sub marking is then a negative loop for p^*, and a *zero-loop* for all $p \neq p^*$, i. e., we can find α and a negative loop τ such that $\alpha' = \alpha \cdot \tau$. In particular, the negative loop τ has a small displacement. The sequence $\alpha \cdot \tau \cdot \beta$ is enabled since we always extend α' by transitions that consume tokens from places where the number of tokens is large (i. e., above the appropriately defined interval). Hence, (i) is satisfied.

Furthermore, since the loop τ only decreases the number of tokens at a place which has a large number of tokens, we can argue that the sequence $\alpha \cdot \beta$ is enabled at $\mu_0 - \Delta(\tau)$ (in other words, the inverse of $\alpha \cdot \beta$ is enabled in the original ff-PN $\mathcal{P}$). Hence, (ii) is satisfied.

We still must argue that (iii) and (iv) are satisfied. In fact, this requires a more carefully procedure. In particular, we must preprocess σ in a certain way to obtain a subsequence which only uses places that are not too small. This subsequence is then processed. Let μ' denote the marking reached by this subsequence. Without going into detail, the main idea is that μ' is large enough such that, in the original ff-PN $\mathcal{P}$, μ' enables each positive loop that is the inverse of some negative loop τ_i created by the permutation procedure. Therefore, (iv) is satisfied.

Actually, (iii) is not necessarily satisfied. The displacement of τ might be small but the length of the sequence itself is unbounded. To solve this problem, we can extract a short sequence from τ that is also a negative loop for p^* and a zero-loop for all other places. Here, we also left out many details in the procedure that ensure that we can find such a short negative loop, and that this negative loop is compatible with (i)–(iv). All these details are discussed later.

6 Canonical Firing Sequences in Fork-Free Petri Nets

In this section, we show that the class of ff-PNs is f-2-canonical for some function f which is exponentially bounded in the size of the input net. In particular, we show that each marking reachable in an ff-PN has a canonical firing sequence with the properties of Definition 4.1

On a high level, we achieve this by providing a procedure that takes some sequence σ leading to a reachable marking μ, and construct from this sequence the sequence ξ. (More precisely, by doing so we show that a sequence satisfying Definition 4.1 for μ exists.) This procedure permutes the given sequence and replaces parts of it with *short* displacement-equivalent sequences.

The stepping stone for the construction is a procedure that, given a firing sequence of a jf-PN (satisfying some conditions), permutes it in such a way that we obtain a firing sequence which, among other things, has a suffix which is a negative loop and satisfies some other useful properties. As already explained in the previous chapter, the reason why we consider jf-PNs instead of ff-PNs is because it is intuitively easier to verify in case of a jf-PN that a some transition is enabled at some marking since each transition only depends on a single place. Furthermore, by considering jf-PNs, we can immediately use many known results for this class of Petri nets.

Lemma 6.1. *Let (P, T, F, μ_0) be a jf-PN, and σ a firing sequence leading to some marking μ_σ such that $\mu_\sigma(p) \geq W$ for all $p \in {}^\bullet\sigma$, and $\mu_0(p^*) \geq \mu_\sigma(p^*) + (\delta + 3) \cdot W$ for some place $p^* \in {}^\bullet\sigma$,*

where $\delta = (\max(\mu_\sigma) + W + 1)^{n-1}$. *Then, there are transition sequences* α, τ *and a marking* $\mu^* \in \{0, W\}^n$ *such that*

(*a*) $\alpha \cdot \tau$ *is a permutation of* σ *with* $\mu_0 \xrightarrow{\alpha\cdot\tau} \mu_\sigma$,

(*b*) $(\mu_0 + \Delta(\tau)) \xrightarrow{\alpha} \mu_\sigma$,

(*c*) τ *is a negative loop with* $\Delta(\tau)(p^*) \in [-(nmW+1)^{n+m}, -1]$ *and* $\Delta(\tau)(p) = 0$ *for all* $p \neq p^*$,

(*d*) $|\tau| \leq (nmW + 1)^{n+m}$, *and*

(*e*) $(\mu^* - \Delta(\tau)) \xrightarrow{\tau} \mu^*$, *where* $\mu^*(p) = W$ *if* $\mu_\sigma(p) \geq W$, *and* $\mu^*(p) = 0$ *otherwise.*

Proof. In the following, we will construct transition sequences α', τ' satisfying (*a*) and (*b*) (replace α and τ by α' and τ' there) as well as the property

(*c'*) $\Delta(\tau)(p^*) \in [-\delta \cdot W, -1]$ and $\Delta(\tau)(p) = 0$ for all $p \in P \setminus \{p^*\}$.

Property (*c'*) is weaker than (*c*) in the sense that (*c*) implies (*c'*) but not the other way round. We can later nevertheless use α' and τ' to define the sequences α and τ of interest which satisfy the properties of the lemma. To find α' and τ', we use a procedure which, by permuting and splitting σ in an appropriate way, creates a number of intermediate sequences which are then recombined to α' and τ'.

Our first step to this end consists of finding transition sequences α_0, φ_0, ..., φ_δ and markings ν_0, ..., ν_δ such that, for appropriately defined integral levels $\ell_0 > \ldots > \ell_{\delta+1}$, the following properties are satisfied for all $i \in [0, \delta]$:

(*i*) $\nu_0 \xrightarrow{\varphi_0 \cdots \varphi_i} \nu_i$,

(*ii*) $\nu_i(p) \in [0, \max(\mu_\sigma) + W]$ for all $p \in P \setminus \{p^*\}$, and

(*iii*) $\nu_i(p^*) \in [\ell_{i+1} + 1, \ell_i]$.

Moreover, these transition sequences shall be defined in such a way that we can use them later to construct α' and τ'.

The important property about these sequences is that, for each $p \neq p^*$, the number of tokens at p is trapped in the interval $[0, \max(\mu_\sigma)+W]$ at each marking ν_i, and that the number of tokens at p^* is strictly monotonically decreasing by a small amount when proceeding from ν_i to ν_{i+1}. This means that some sub-marking restricted to the places $p \neq p^*$ must repeat at some point (i. e., $\nu_i(p) = \nu_j(p)$ for some i and j and all $p \neq p^*$), and hence we have found a negative loop (with displacement $\nu_j - \nu_i$). The general idea how we find these markings and sequences is illustrated in Figure 1.

The role of the marking ν_0 is to serve as an alternative initial marking. This marking is the starting point for constructing the sequences above. The important property of ν_0 is that, for each place $p \neq p^*$, the absolute value of the difference between $\nu_0(p)$ and $\mu_\sigma(p)$ is small (remember that we want to trap $p \neq p^*$ in the interval $[0, \max(\mu_\sigma) + W]$).

We first show how to construct ν_0. We initialize $\alpha_0 \leftarrow \epsilon$ as the empty sequence and $\psi \leftarrow \sigma$. As long as there is a $t \in \Psi_{\text{first}}(\psi)$ such that ${}^\bullet t \neq p^*$ and $(\mu_0 + \Delta(\alpha_0))({}^\bullet t) > \mu_\sigma({}^\bullet t) + W$, we set $\alpha_0 \leftarrow \alpha_0 \cdot t$ and $\psi \leftarrow \psi \dot{-} t$. (Note that this means that we extend α_0 by firing t at marking $\mu_0 + \Delta(\alpha_0)$.) We stop, when no such t exists (any more). After that, we define $\nu_0 := \mu_0 + \Delta(\alpha_0)$ and $\varphi_0 := \epsilon$. (The sequence φ_0 is a dummy sequence to avoid a special case when constructing the remaining sequences φ_i, $i > 0$.) By Lemma 3.2, we observe $\mu_0 \xrightarrow{\alpha_0} \nu_0 \xrightarrow{\psi} \mu_\sigma$.

We remark that, in the procedure above, there may be different possibilities to choose $t \in \Psi_{\text{first}}(\psi)$. Depending on which of them we choose, the resulting sequences may be different. If we want to make them unique, we could define an ordering of the transitions and then choose the smallest possible transition w. r. t. this ordering. However, to show the lemma, it is not important which transition we choose. This remark also applies at various occasions later when we permute sequences using similar schemes.

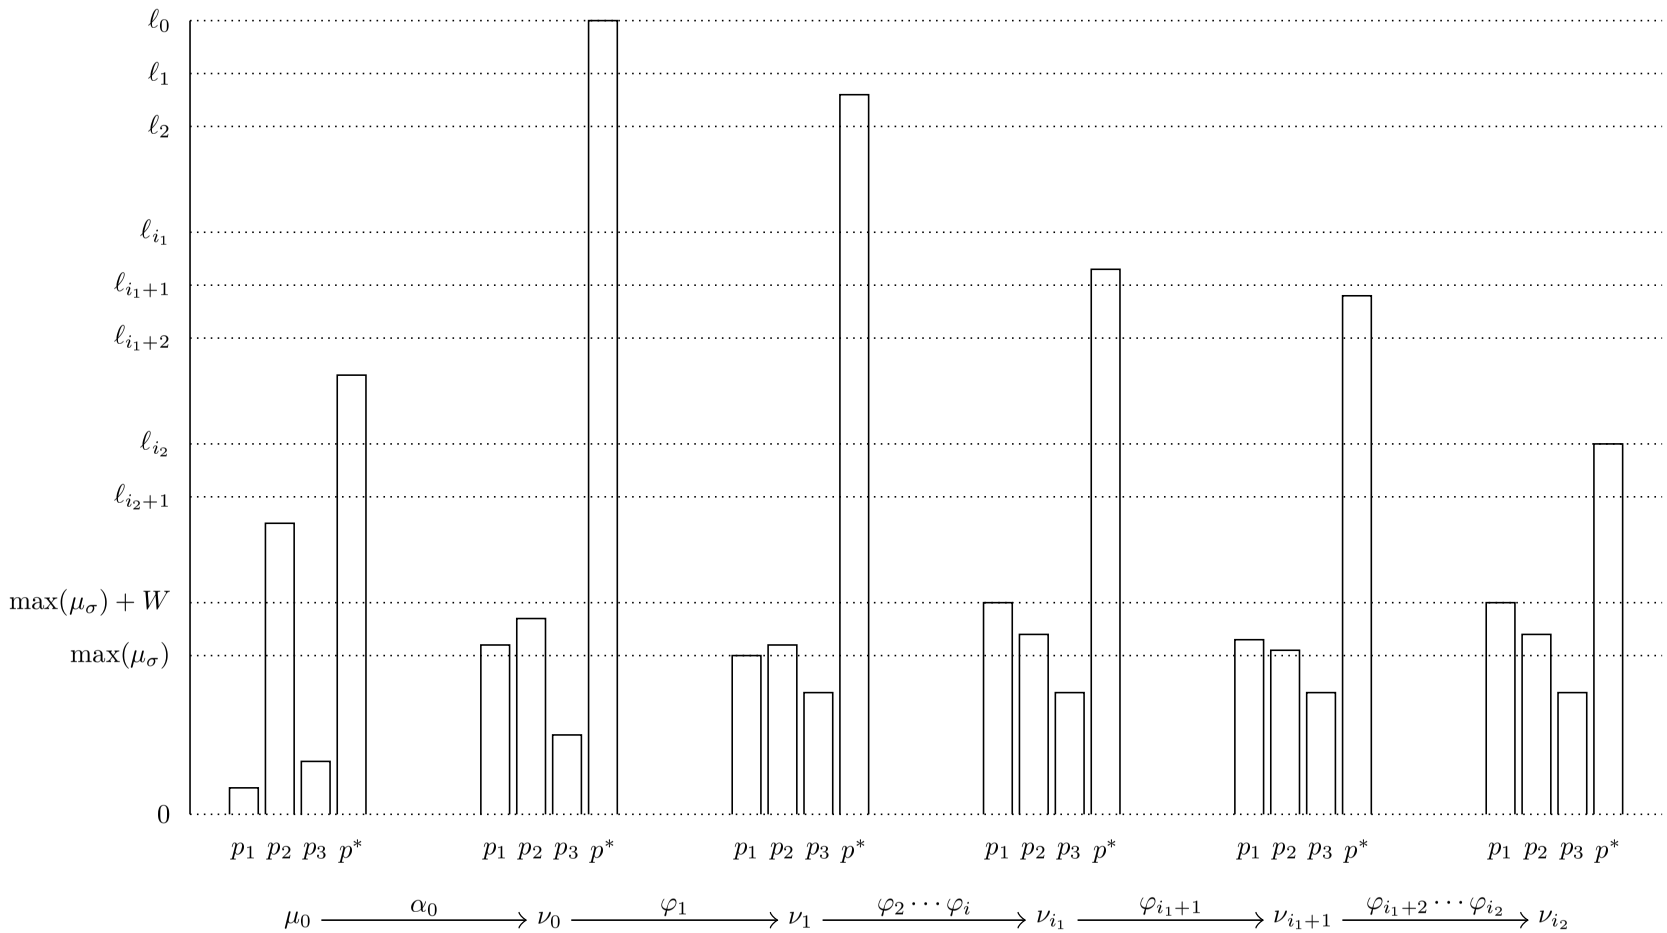

Figure 1: The sequences generated in the following consist of transitions of the original sequence σ such that each occurrence of a transition of σ is used at most once in these sequences. Starting from μ_0, we generate α_0 which leads to a marking ν_0 such that $\nu_0(p) \leq \max(\mu_\sigma)(p) + W$ for all $p \in P \setminus \{p^*\}$, and $\Delta(\alpha_0)(p^*) \geq 0$. From there, we define levels ℓ_i, $i \in [0, \delta + 1]$, and generate sequences φ_i, $i \in [\delta]$, such that, when encountering the respective corresponding markings ν_i, the respective number of tokens at each place $p \in P \setminus \{p^*\}$ is still trapped within the interval $[0, \max(\mu_\sigma) + W]$, and such that the number of tokens at p^* is strictly decreasing, where $\nu_i(p^*) \in [\ell_{i+1} + 1, \ell_i]$. By constructing a sufficient number of these sequences and markings, we find $i_1 < i_2$ such that $\nu_{i_1}(p) = \nu_{i_2}(p)$ for all $p \in P \setminus \{p^*\}$. The sequence $\varphi_{i_1+1} \cdots \varphi_{i_2} =: \tau''$ is then a negative loop for p^*. Permuting τ'' in an appropriate way yields the negative loop τ'.

Next, we define the levels $\ell_0, \ldots, \ell_{\delta+1}$ by $\ell_i := \nu_0(p^*) - i \cdot W$, $i \in [0, \delta+1]$. Note that properties (i)–(iii) are satisfied for $i = 0$.

We proceed by recursively defining the remaining transition sequences and markings. For all $i \in [\delta]$, let φ_i and ν_i be recursively defined by $\varphi_i := \text{getSubSeq}(i)$ and $\nu_i := \nu_0 + \Delta(\varphi_0 \cdots \varphi_i)$. Furthermore, we define $\sigma_{\text{rest}} := \sigma \dot{-} (\alpha_0 \cdot \varphi_0 \cdots \varphi_\delta)$. The function getSubSeq basically does the

Function getSubSeq(iteration i)

$\varphi \leftarrow$ empty sequence
$\psi_{\text{rest}} \leftarrow \psi \dot{-} (\varphi_0 \cdots \varphi_{i-1})$
while $\exists t \in \Psi_{first}(\psi_{rest}) : [(p^* = {}^\bullet t$ **and** $(\nu_{i-1} + \Delta(\varphi))(p^*) > \ell_i)$ **or**
$(\exists p \in P \setminus \{p^*\} : p = {}^\bullet t$ **and** $(\nu_{i-1} + \Delta(\varphi))(p) > \mu_\sigma(p) + W)]$ **do**
  $\varphi \leftarrow \varphi \cdot t$
  $\psi_{\text{rest}} \leftarrow \psi_{\text{rest}} \dot{-} t$
return φ

following: Starting from the marking ν_{i-1}, it creates a sequence φ_i by iteratively choosing transitions from the sequence containing the remaining transitions (which is $\psi \dot{-} (\varphi_0 \cdots \varphi_{i-1})$ at the beginning of the function). At each step, it either chooses a transition that consumes tokens from p^* provided p^* has too many tokens (i. e., the token number is larger than level ℓ_i) or chooses a transition that consumes tokens from some place $p \neq p^*$ provided p has too many tokens (i. e., the token number is larger than $\mu_\sigma(p) + W$). Note that it is always possible to find such a transition since, otherwise, the remaining sequence would not lead to μ_σ.

Lemma 3.2 immediately implies that $\alpha_0 \cdot \varphi_0 \cdots \varphi_\delta \cdot \sigma_{\text{rest}}$ is a firing sequence which ensures property (i) for all $i \in [\delta]$. Using induction on $i \in [0, \delta]$, it is not hard to show properties (ii)–(iii) for the remaining sequences and markings: Let $i \in [\delta]$, and assume that (ii)–(iii) hold for step $i-1$. Property (ii) and $\nu_i(p^*) \leq \ell_i$ of (iii) directly follow from the definition of Function getSubSeq. Furthermore, $\nu_{i-1}(p^*) \geq \ell_i + 1$ holds by the induction hypothesis. Subsequently, one iteration of the while-loop of Function getSubSeq cannot decrease the number of tokens at p^* below $\ell_{i+1} + 1$, which implies $\nu_{i+1}(p^*) \geq \ell_{i+1} + 1$. Thus, (iii) holds.

We continue the proof by defining α' and τ'. By (ii), there are at most δ different projections of the $\delta + 1$ markings ν_i, $i \in [0, \delta]$, onto the places $p \in P \setminus \{p^*\}$. Hence, there are i_1, $i_2 \in [0, \delta]$, $i_1 < i_2$, such that $\nu_{i_1}(p) = \nu_{i_2}(p)$ for all $p \in P \setminus \{p^*\}$. Let $\alpha_1 := \varphi_0 \cdots \varphi_{i_1}$, $\tau'' := \varphi_{i_1+1} \cdots \varphi_{i_2}$, $\alpha_2 := \varphi_{i_2+1} \cdots \varphi_\delta \cdot \sigma_{\text{rest}}$, $\alpha_{0,1} := \alpha_0 \cdot \alpha_1$, $\alpha_{1,2} := \alpha_1 \cdot \alpha_2$, and $\alpha' := \alpha_0 \cdot \alpha_1 \cdot \alpha_2$, see (a) of Figure 2. (In this illustration, certain sequences are indicated to be enabled at certain markings. This follows from the definition of the sequences and from our observations made above, in particular property (i).)

Property (iii) ensures $\Delta(\tau'')(p^*) \in [-(i_2 - i_1 + 1)W, -1]$, and therefore that τ'' satisfies (c') (where τ of (c') is identified with τ''). This together with the assumption of this lemma implies $(\nu_0 + \Delta(\alpha_{1,2}))(p) \geq \mu_\sigma(p) \geq W$ for all $p \in {}^\bullet\tau''$. Therefore, we can apply Lemma 3.3 to the markings $\nu_0 + \Delta(\alpha_{1,2})$ (which is none other than $\mu_0 + \Delta(\alpha')$), μ_σ, and the sequence τ'' to obtain a permutation τ' of τ'' which is enabled at $\mu_0 + \Delta(\alpha')$. Note that also τ' satisfies (c') (where τ of (c') is identified with τ'). Next we observe that α_2 is enabled at ν_{i_1} since α_2 is enabled at ν_{i_2} and $\nu_{i_1} \geq \nu_2$ holds. Consequently, $\alpha' \cdot \tau'$ is a permutation of σ enabled at μ_0, and (a) follows, see (b) of Figure 2.

To show (b) for α', it's only necessary to consider $\alpha_{0,1}$ since we already know that $\nu_{i_2} \xrightarrow{\alpha_2} \mu_\sigma$. Moreover, it's sufficient to verify that, for all $j \in [0, |\alpha_{0,1}| - 1]$ and $p \in P$, $(\mu_0 + \Delta(\tau') + \Delta((\alpha_{0,1})_{[..j]}))(p) \geq \min\{W, (\mu_0 + \Delta((\alpha_{0,1})_{[..j]}))(p)\}$ holds since this implies that the marking $(\mu_0 + \Delta(\tau') + \Delta((\alpha_{0,1})_{[..j]}))$ enables the next transition $(\alpha_{0,1})_{[j+1]}$. Since, by (c'), τ' does not change the number of tokens at some place other than p^*, this sufficient condition remains to be shown for p^*, which we do now.

We first consider α_0. Since $\mu_0(p^*) + \Delta(\tau')(p^*) \geq W$ and $p^* \notin {}^\bullet\alpha_0$, we find $(\mu_0 + \Delta(\tau') + \Delta((\alpha_0)_{[..j]}))(p^*) \geq W$ for all $j \in [0, |\alpha_0| - 1]$.

Next, we consider $\alpha_1 = \varphi_1 \cdots \varphi_{i_1}$. By (iii), we find

$$\begin{aligned}\nu_i(p^*) &\geq \ell_{i+1} + 1 > \ell_{i_1} = \ell_{i_1+1} + W = \nu_0(p^*) - (i_1+1)\cdot W + W \\ &\geq \mu_0(p^*) - (i_1+1)\cdot W + W \geq \mu_\sigma(p^*) + (\delta+3)\cdot W - (i_1+1)\cdot W + W \\ &= \mu_\sigma(p^*) + (\delta - i_1 + 1)\cdot W + 2W \geq (\delta - i_1 + 1)\cdot W + 3W\end{aligned}$$

for all $i \in [0, i_1 - 1]$ (remember for this sequence of inequalities the prerequisite $\mu_\sigma(p^*) \geq W$). By property (iii) and the definition of Function getSubSeq, we observe $\Delta((\varphi_i)_{[..j]})(p^*) \geq -2W$ for all $i \in [0, \delta]$, $j \in [0, |\varphi_i|]$. Thus, $(\nu_i + \Delta((\varphi_{i+1})_{[..j]}))(p^*) \geq (\delta - i_1 + 1)\cdot W + W$ for all $i \in [0, i_1 - 1]$ and $j \in [0, |\varphi_{i+1}|]$. Adding the displacement of τ' yields $(\nu_i + \Delta(\tau') + \Delta((\varphi_{i+1})_{[..j]}))(p^*) \geq -(i_2 - i_1 + 1)W + (\delta - i_1 + 1)\cdot W + W \geq W$ for all $i \in [0, i_1 - 1]$ and $j \in [0, |\varphi_{i+1}|]$. This implies $(\nu_0 + \Delta(\tau') + \Delta((\alpha_1)_{[..j]}))(p^*) \geq W$ for all $j \in [0, |\alpha_1| - 1]$. In total, (b) follows for α', see (c) of Figure 2.

So far, we have shown that α' and τ' satisfy (a), (b), and (c'), but not necessarily (c). We now define sequences α and τ satisfying (a)–(c). The following observations are illustrated in

(a)

$$\mu_0 \xrightarrow{\alpha_0} \nu_0 \xrightarrow{\alpha_1 = \varphi_0 \cdots \varphi_{i_1}} \nu_{i_1} \xrightarrow{\tau'' = \varphi_{i_1+1} \cdots \varphi_{i_2}} \nu_{i_2} \xrightarrow{\varphi_{i_2+1} \cdots \varphi_\delta} \nu_\delta \xrightarrow{\sigma_{\text{rest}}} \mu_\sigma$$

$$\nu_{i_2} \xrightarrow{\alpha_2 = \varphi_{i_2+1} \cdots \varphi_\delta \cdot \sigma_{\text{rest}}} \mu_\sigma$$

(b)

$$\mu_0 \xrightarrow{\alpha' = \alpha_0 \cdot \alpha_1 \cdot \alpha_2} \mu_0 + \Delta(\alpha') \xrightarrow{\tau'} \mu_\sigma$$

(c)

$$\mu_0 + \Delta(\tau') \xrightarrow{\alpha' = \alpha_0 \cdot \alpha_1 \cdot \alpha_2} \mu_\sigma$$

(d)

$$\forall i \in [q]: \ \mu_\sigma - \Delta(\tau_i) \xrightarrow{\tau_i} \mu_\sigma$$

$$\mu_\sigma - \Delta(\alpha) = \mu_0 + \Delta(\tau_j) \xrightarrow{\alpha = \alpha_0 \cdot \alpha_1 \cdot \alpha_2 \cdot \tau_1 \cdots \tau_{j-1} \cdot \tau_{j+1} \cdots \tau_q} \mu_\sigma$$

$$\mu_0 \xrightarrow{\alpha = \alpha_0 \cdot \alpha_1 \cdot \alpha_2 \cdot \tau_1 \cdots \tau_{j-1} \cdot \tau_{j+1} \cdots \tau_q} \mu_0 + \Delta(\alpha) \xrightarrow{\tau = \tau_j} \mu_\sigma$$

Figure 2: Observations for different steps of Lemma 6.1

(d) of Figure 2. We apply Lemma 3.4 to the inverse net of $\mathcal{P}$ and $\Psi(\tau')$, and obtain nonpositive loops $\Phi_1, \ldots, \Phi_q$ with $\min(\Phi_i) \geq -(1+(n+m)W)^{n+m} \geq -(nmW+1)^{c(n+m)}$ for all $i \in \ell$ and some constant $c \in \mathbb{N}_{>0}$, and $\Psi(\tau') = \sum_{i=1}^{q} \Phi_i$. By property (c') and by the fact that $(\mu_\sigma - \Delta(\Phi_i))(p) \geq \mu_\sigma(p) \geq W$ holds for all $p \in {}^\bullet\sigma \supseteq {}^\bullet\tau_i$ and $i \in [q]$, we can apply Lemma 3.3 to find nonpositive loops $\tau_1, \ldots, \tau_q$ with $\Psi(\tau_i) = \Phi_i$ and $\mu_\sigma - \Delta(\tau_i) \xrightarrow{\tau_i} \mu_\sigma$ for all $i \in [q]$ such that, for some $j \in [q]$, $\tau_j =: \tau$ is a negative loop satisfying (c) and (d). Correspondingly, we define $\alpha := \alpha' \cdot \tau_1 \cdots \tau_{j-1} \cdot \tau_{j+1} \cdots \tau_q$. Note that $\alpha \cdot \tau$ is a firing sequence since α' is enabled at $\mu_\sigma - \Delta(\alpha')$ (i. e., property (b)) and each τ_i, $i \in [q]$, is a nonpositive loop enabled at $\mu_\sigma - \Delta(\tau_i)$. Hence, α and τ satisfy (a)–(e). □

We obtain the following canonical firing sequence by first permuting a given firing sequence appropriately, and then iteratively applying Lemma 6.1 as well as Lemma 3.5. Since the most characteristic property of this canonical sequence is that almost all transitions are contained in short negative loops, we call it *negative canonical sequence*.

Lemma 6.2. *There is a constant $c \in \mathbb{N}_{>0}$ such that, for each jf-PN $\mathcal{P} = (N, \mu_0)$ with $n > 0$ and each marking μ reachable in $\mathcal{P}$, there are $k \in [0, n \cdot \max(\mu_0)]$, transition sequences α, β, $\tau_1, \ldots, \tau_k$, and a marking $\mu^* \in \{0, W\}^n$ such that*

- (a) *$\alpha \cdot \tau_1 \cdots \tau_k \cdot \beta$ is a firing sequence leading from μ_0 to μ,*
- (b) *$\alpha \cdot \beta$ is enabled at $(\mu - \Delta(\alpha \cdot \beta))$ with $|\alpha \cdot \beta| \leq (nmW + \max(\mu) + 1)^{cn^2(n+m)}$,*
- (c) *$\mu_0 + \Delta(\alpha \cdot \tau_1 \cdots \tau_k) \geq \mu^*$, and*
- (d) *each τ_i, $i \in [k]$, is a negative loop with*
  - (i) *$\Delta(\tau_i)(p^*) \in [-(nmW+1)^{(n+m)}, -1]$ for some place p^* and $\Delta(\tau_i)(p) = 0$ for all $p \neq p^*$,*
  - (ii) *$|\tau_i| \leq (nmW+1)^{(n+m)}$, and*
  - (iii) *$(\mu^* - \Delta(\tau_i)) \xrightarrow{\tau_i} \mu^*$.*

Proof. We assume m, $W > 0$ since the lemma trivially holds otherwise. Let μ be a reachable marking, and ρ be a firing sequence leading to μ.

We want to show the lemma by repeatedly applying Lemma 6.1. However, this lemma requires that the sequence σ which we are applying the lemma to satisfies $\mu_\sigma(p) \geq W$ for all $p \in {}^\bullet\sigma$. Therefore, we extract such a sequence from ρ as follows.

First, we initialize $\sigma \leftarrow \epsilon$ and $\rho_{\text{rest}} \leftarrow \rho$. As long as there is a transition $t \in \Psi_{\text{first}}(\rho_{\text{rest}})$ with $(\mu_0 + \Delta(\sigma))({}^\bullet t) \geq 2W$, we set $\sigma \leftarrow \sigma \cdot t$ and $\rho_{\text{rest}} \leftarrow \rho_{\text{rest}} \dot{-} t$. At the end of this procedure, we define $\mu_\sigma := \mu_0 + \Delta(\sigma)$. It is easy to see that σ has the desired property.

Now, we iteratively apply Lemma 6.1 as long as possible, using $(P, T, F, \mu_0 + \Delta(\tau_1 \cdots \tau_{i-1}))$, α_{i-1}, and μ_σ in the i-th iteration, where τ_i and α_i (with $\alpha_0 = \sigma$) denote the sequences resulting from the i-th application of the lemma. Let $\ell \leq n \cdot \max(\mu_0)$ denote the number of applications of Lemma 6.1, and μ^* the marking defined in Lemma 6.1. We remark that α_ℓ is the "residual" sequence which remains of σ after we have extracted the negative loops $\tau_1, \ldots, \tau_\ell$. We first observe that, by Lemma 6.1, the negative loops τ_i, $i \in [\ell]$, satisfy property (d).

Next, we find $(\mu_\sigma - \Delta(\alpha_\ell)) \xrightarrow{\alpha_\ell} \mu_\sigma \geq \mu^*$. Furthermore, $\max(\mu_\sigma) \leq \max(\mu) + 2W$ holds. To see this, assume $\max(\mu_\sigma) > \max(\mu) + 2W$ for the sake of contradiction. Then, there is a transition $t \in \Psi_{\text{first}}(\rho_{\text{rest}})$ with $(\mu_0 + \Delta(\sigma))({}^\bullet t) \geq 2W$ which, by the construction of σ, cannot be the case. By this, we find

$$\begin{aligned}
&\max(\mu_\sigma - \Delta(\alpha_\ell)) \leq \max(\mu_\sigma) + (\delta + 3) \cdot W \\
&\leq \max(\mu_\sigma) + ((\max(\mu_\sigma) + W + 1)^{n-1} + 3) \cdot W \\
&\leq (\max(\mu) + 2W)^{an}
\end{aligned}$$

for some constant $a \in \mathbb{N}_{>0}$. We apply Lemma 3.5 to the markings $(\mu_\sigma - \Delta(\alpha_\ell))$ and μ_σ, and obtain a transition sequence α (i. e., sequence ξ of Lemma 3.5) with $(\mu_\sigma - \Delta(\alpha_\ell)) = (\mu_\sigma - \Delta(\alpha)) \xrightarrow{\alpha} \mu_\sigma$. The sequence α is not necessarily a permutation of α_ℓ but has the same displacement. In particular, we have $\mu_0 \xrightarrow{\alpha} (\mu_0 + \Delta(\alpha)) \xrightarrow{\tau_1 \cdots \tau_\ell} \mu_\sigma \geq \mu^*$, meaning that (c) is satisfied. Let c denote the constant of Lemma 3.5 (the constant c of the Lemma we are proving at the moment is somewhat larger than the constant c of Lemma 3.5). Furthermore, let u and k be the numbers of this particular application of Lemma 3.5. Using the bounds given above, we find

$$\begin{aligned}
|\alpha| \leq (k+1)u &\leq \Big(\max(\mu_\sigma) + (nmW + \max(\mu_\sigma - \Delta(\alpha_\ell)) + 1)^{cn(n+m)}\Big)^2 \\
&\leq \Big((\max(\mu) + 2W) + (nmW + (\max(\mu) + 2W)^{an} + 1)^{cn(n+m)}\Big)^2 \\
&\leq (nmW + \max(\mu) + 1)^{dn^2(n+m)}
\end{aligned}$$

for some constant $d \in \mathbb{N}_{>0}$. Last, we apply Lemma 3.5 to the markings μ_σ and μ, yielding a transition sequence β (i. e., sequence ξ of Lemma 3.5) with $\mu_\sigma \xrightarrow{\beta} \mu$. Let u and k denote the numbers of this particular application of Lemma 3.5. Then, we observe

$$\begin{aligned}
|\beta| \leq (k+1)u &\leq \Big(\max(\mu) + (nmW + \max(\mu_\sigma) + 1)^{cn(n+m)}\Big)^2 \\
&\leq \Big(\max(\mu) + (nmW + \max(\mu) + 2W + 1)^{cn(n+m)}\Big)^2 \\
&\leq (nmW + \max(\mu) + 1)^{dn(n+m)}
\end{aligned}$$

for some constant $d \in \mathbb{N}_{>0}$. The observations above imply $(\mu_\sigma - \Delta(\alpha)) \xrightarrow{\alpha} \mu_\sigma \xrightarrow{\beta} \mu$ as well as $\mu_0 \xrightarrow{\alpha} (\mu_0 + \Delta(\alpha)) \xrightarrow{\tau_1 \cdots \tau_k} \mu_\sigma \xrightarrow{\beta} \mu$. This together with the bounds on $|\alpha|$ and $|\beta|$ ensures properties (a) and (b). □

By reversing negative canonical firing sequences of jf-PNs, we obtain positive canonical firing sequences of ff-PNs.

Corollary 6.3. *There is a constant $c \in \mathbb{N}_{>0}$ such that, for each ff-PN $\mathcal{P} = (N, \mu_0)$ with $n > 0$ and each marking μ reachable in $\mathcal{P}$, there are $k \in [0, n \cdot \max(\mu)]$, transition sequences α, β, τ_1, …, τ_k, and a marking $\mu^* \in \{0, W\}^n$ such that*

- (a) *$\alpha \cdot \tau_1 \cdots \tau_k \cdot \beta$ is a firing sequence leading from μ_0 to μ,*
- (b) *$\alpha \cdot \beta$ is enabled at μ_0 with $|\alpha \cdot \beta| \leq (nmW + \max(\mu_0) + 1)^{cn^2(n+m)}$,*
- (c) *$\mu_0 + \Delta(\alpha) \geq \mu^*$, and*
- (d) *each τ_i, $i \in [k]$, is a positive loop with*
  - (i) *$\Delta(\tau_i)(p^*) \in [1, (nmW + 1)^{c(n+m)}]$ for some place p^* and $\Delta(\tau_i)(p) = 0$ for all $p \neq p^*$,*
  - (ii) *$|\tau_i| \leq (nmW + 1)^{c(n+m)}$, and*
  - (iii) *τ_i is enabled at μ^*.*

Proof. Lemma 6.2 implies the corollary by considering the inverse net of $\mathcal{P}$, as well as μ as the initial marking and μ_0 as the end marking. □

Lemma 6.4. *There is a constant $c \in \mathbb{N}_{>0}$ such that the class of ff-PNs is simple structurally f-g-canonical, where $f(n, m, W, K) = (nmW + K + 1)^{cn^2(n+m)}$ and $g(n, m, W, K) = 2$.*

Proof. This follows immediately from Corollary 6.3. □

7 Complexity Results

Our observations of the previous section allows us to apply the framework to obtain the following complexity results.

Theorem 7.1. *The RecLFS, (zero-)reachability, boundedness, and covering problems of ff-PNs are* **PSPACE**-*complete, even if restricted to WSMs. The liveness problem of ff-PNs is in* **PSPACE**.

Proof. **PSPACE**-completeness for the RecLFS problem is immediately implied by Lemma 3.1 (for an Petri net $\mathcal{P} = (P, T, F, \mu_0)$ and Parikh vector Φ consider the inverse net of $\mathcal{P}$ with $\mu_0 + \Delta(\Phi)$ as initial marking). The lower bound for the (zero-)reachability, boundedness, and covering problems are provided by Lemma 3.6. By Lemma 6.4, the class of ff-PNs is simple structurally f-f-canonical for the function f given in the lemma, where $f(n, m, W, K) = (nmW + K + 1)^{cn^2(n+m)}$. By Lemma 4.2, the (zero-)reachability, and covering problems of ff-PNs are decidable in space polynomial in $\text{size}(\mathcal{P}) + \text{size}(\mu) + n \operatorname{ld} f(n, m, W, \max(\mu_0)) + r \overset{P}{\leq} \text{size}(\mathcal{P}) + \text{size}(\mu)$. Similarly, the boundedness problem is decidable in space polynomial in $\text{size}(\mathcal{P}) + n \operatorname{ld} f(n, m, W, \max(\mu_0)) + r \overset{P}{\leq} \text{size}(\mathcal{P})$. By Lemma 4.3, the liveness problem of ff-PNs is decidable in space polynomial in $\text{size}(\mathcal{P}) + n \operatorname{ld}(f(n, m, W, \max(\mu_0) + f(n, m, W, \max(\mu_0))W)) + r \overset{P}{\leq} \text{size}(\mathcal{P})$. □

Theorem 7.2. *The containment and the equivalence problems of ff-PNs are decidable in exponential space and are* **PSPACE**-*hard, even if restricted to WSMs.*

Proof. **PSPACE**-hardness follows from Lemma 3.7. Let s denote the encoding size of the input. Then, by Lemma 4.4 and Lemma 6.4, there are polynomials p, p' such that the equivalence and containment problems are decidable in space polynomial in $(K + 2u_1K)^{u_2 n} \cdot 2^{p(s+n \operatorname{ld}(u_3))} + r \overset{P}{\leq} (K + 2f(n, m, W, K)K)^{2n} \cdot 2^{p(s+n \operatorname{ld} f(n, m, W, K+2f(n, m, W, K)W))} \overset{P}{\leq} 2^{p'(s)}$. □

8 Conclusion

We investigated fork-free Petri nets, a natural class of Petri nets. Using results for join-free Petri nets, we showed that reachable markings in ff-PNs can be reached by firing sequences with nice properties. This enabled us to apply a mathematical framework for Petri nets, which provided upper bounds for many classical computational problems. We showed that the (zero-)reachability, boundedness, and covering problems are **PSPACE**-complete, that the liveness problem is in **PSPACE**, and that the containment and equivalence problems are **PSPACE**-hard and decidable in exponential space. Clearly, for the latter three problems, the respective lower and upper bounds do not match. Further research is needed to fill this gap. Moreover, it would be interesting to know if there are other classes for which completeness-results for the classical problems investigated in this paper are unknown and for which the framework of [6] can be applied.

References

1. C. Amer-Yahia, N. Zerhouni, A. E. Moudni, and M. Ferney. Some subclasses of Petri nets and the analysis of their structural properties: a new approach. *IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans*, 29(2):164–172, 1999.

2. A. W. Holt and F. Commoner. Events and conditions. In *Record of the Project MAC Conference on Concurrent Systems and Parallel Computation*, pages 3–52. ACM, New York, 1970.

3. R. M. Karp and R. E. Miller. Parallel program schemata. *Journal of Computer and System Sciences*, 3(2):147–195, 1969.

4. Y. E. Lien. Termination properties of generalized Petri nets. *SIAM Journal on Computing*, 5 (2):251–265, 1976.

5. E. W. Mayr and J. Weihmann. Completeness results for generalized communication-free Petri nets with arbitrary edge multiplicities. In *Proceedings of the 7th International Workshop on Reachability Problems*, volume 8169 of *Lecture Notes in Computer Science*, pages 209–221. Springer, 2013.

6. E. W. Mayr and J. Weihmann. A framework for classical Petri net problems: Conservative Petri nets as an application. In *Proceedings of the 35th International Conference on Application and Theory of Petri Nets (ICATPN'14)*, volume 8489 of *Lecture Notes in Computer Science*, pages 314–333. Springer, 2014.

7. E. W. Mayr and J. Weihmann. Completeness results for generalized communication-free Petri nets with arbitrary arc multiplicities. *Fundamenta Informaticae*, 2015. To appear.

8. K. Morita and T. Watanabe. The legal firing sequence problem of Petri nets with state machine structure. In *Proceedings of the 1996 IEEE International Symposium on Circuits and Systems (ISCAS'96)*, volume 3, pages 64–67, 1996.

9. S. Taoka and T. Watanabe. A linear time algorithm solving the legal firing sequence problem for a class of edge-weighted cactuses. In *Proceedings of the 1999 IEEE International Conference on Systems, Man, and Cybernetics (SMC'99)*, volume 3, pages 893–898, 1999.

10. E. Teruel and M. Silva. Liveness and home states in equal conflict systems. In *Proceedings of the 14th International Conference on Application and Theory of Petri Nets (ICATPN'93)*, volume 691 of *Lecture Notes in Computer Science*, pages 415–432. Springer, 1993.

11. E. Teruel and M. Silva. Well-formedness of equal conflict systems. In *Proceedings of the 15th International Conference on Application and Theory of Petri Nets (ICATPN'94)*, volume 815 of *Lecture Notes in Computer Science*, pages 491–510. Springer, 1994.

12. E. Teruel and M. Silva. Structure theory of equal conflict systems. *Theoretical Computer Science*, 153:271–300, 1996.